

RATE GROUPING DURING CODE GENERATION FOR MULTI-RATE MODELS

Technical Field of the Invention

5 The present invention generally relates to block diagrams and, more particularly, to methods for executing components of a block diagram executing at different rates.

Background

 Various classes of block diagrams describe computations that can be
10 performed on application specific computational hardware, such as a computer, microcontroller, FPGA, and custom hardware. Classes of such block diagrams include time-based block diagrams such as those found within Simulink® from the MathWorks, Inc. Natick Ma, state-based and flow diagrams such as those found within Stateflow® from the MathWorks, Inc. Natick Ma., and data-flow diagrams. A
15 common characteristic among these various forms of block diagrams is that they define semantics on how to execute the diagram.

 Historically, engineers and scientists have utilized time-based block diagram models in numerous scientific areas such as Feedback Control Theory and Signal Processing to study, design, debug, and refine dynamic systems. Dynamic systems,
20 which are characterized by the fact that their behaviors change over time, are representative of many real-world systems. Time-based block diagram modeling has become particularly attractive over the last few years with the advent of software packages such as Simulink from The MathWorks, Inc. of Natick, Massachusetts. Such packages provide sophisticated software platforms with a rich suite of support
25 tools that makes the analysis and design of dynamic systems efficient, methodical, and cost-effective.

 A dynamic system (either natural or man-made) is a system whose response at any given time is a function of its input stimuli, its current state, and the current time. Such systems range from simple to highly complex systems. Physical dynamic
30 systems include a falling body, the rotation of the earth, bio-mechanical systems (muscles, joints, etc.), bio-chemical systems (gene expression, protein pathways), weather and climate pattern systems, etc. Examples of man-made or engineered dynamic systems include: a bouncing ball, a spring with a mass tied on an end,

automobiles, airplanes, control systems in major appliances, communication networks, audio signal processing, nuclear reactors, a stock market, etc. Professionals from diverse areas such as engineering, science, education, and economics build mathematical models of dynamic systems in order to better understand system behavior as it changes with the progression of time. The mathematical models aid in building “better” systems, where “better” may be defined in terms of a variety of performance measures such as quality, time-to-market, cost, speed, size, power consumption, robustness, etc. The mathematical models also aid in analyzing, debugging and repairing existing systems (be it the human body or the anti-lock braking system in a car). The models may also serve an educational purpose of educating others on the basic principles governing physical systems. The models and results are often used as a scientific communication medium between humans. The term “model-based design” is used to refer to the use of block diagram models in the development, analysis, and validation of dynamic systems.

Dynamic systems are typically modeled as sets of differential, difference, and/or algebraic equations. At any given instant of time, these equations may be viewed as relationships between the system’s output response (“outputs”), the system’s input stimuli (“inputs”) at that time, the current state of the system, the system parameters, and time. The state of the system may be thought of as a numerical representation of the dynamically changing configuration of the system. For instance, in a physical system modeling a simple pendulum, the state may be viewed as the current position and velocity of the pendulum. Similarly, a signal-processing system that filters a signal would maintain a set of previous inputs as the state. The system parameters are the numerical representation of the static (unchanging) configuration of the system and may be viewed as constant coefficients in the system’s equations. For the pendulum example, a parameter is the length of pendulum and for the filter example; a parameter is the values of the filter taps.

There are four common types of mathematical models used in the study of dynamic systems. The first type of mathematical model describes systems using ordinary differential equations (ODEs) and is depicted in Figure 1A. The dynamic system 2 specifies a set of two equations: Output 4 and Derivative 6. The Output equation 4 facilitates the computation of the system’s output response at a given time instant as a function of its inputs, states, parameters, and time. The Derivative equation 6 is an ordinary differential equation that allows the computation of the

derivative of the states at the current time as a function of the inputs, the states, parameters, and time. This class of models is suitable for systems in which it is important to track the system response as a continuous function of time. Such continuous-time systems are commonly representative of physical systems (mechanical, thermal, electrical). For simple systems, it may be possible to use the Output 4 and Derivative equations 6 to obtain a closed-form solution for the output response $y(t)$. But in most complex real world systems, the response of the system is obtained by integrating the states through numerical means.

The definition of an ODE used herein encompasses both implicit and explicit differential equations. The class of ordinary differential equations may require additional equations to define the system being modeled. For example, equations called projections may be required to impose constraints on the differential variables (e.g., states X_1 and X_2 must fall on the manifold defined by $X_1^2 + X_2^2 = 25$). These constraints can be either applied as a secondary condition or a coupled condition to the differential equation. Although systems including the projections may conventionally no longer qualify as an ODE; they are included here to simplify the categories of systems. Another example is the use of a Jacobian equation that defines partial derivatives with respect to the independent and/or differential variables. The Jacobian equation is typically used when obtaining a linear approximation of a non-linear model or an overall linear model of a set of equations. Jacobian equations are required for some forms of numerical integration, for producing the linear model once the model has reached its steady state operating point, etc. The Output 4 and Derivatives equations 6 may be extended to define other relationships for the block. For example, the Output equation 4 may help manage its states by defining a relationship where it resets the state back to a known quantity at a specific point in time or when a specific condition is seen.

Another type of mathematical model describes systems using difference equations as depicted in Figure 1B. The dynamic system 8 specifies a set of two equations: Output 10 and Update 12. The Output equation 10 facilitates the computation of the system's output response at a given time instant as a function of the inputs, states at some previous time, parameters, and time. The Update equation 12 is a difference equation that allows the computation of the states at the current time as a function of the inputs, states at some previous time, parameters, and time. This

class of models is suitable for systems in which it is important to track the system response at discrete points in time. Such discrete-time systems are commonly representative of discrete-time control and digital signal processing systems. For simple systems, it may be possible to use the Output 10 and Update equations 12 to obtain a closed-form solution for the output response $y(t)$. But in most complex real world systems, the response of the system is solved through recursion. The Output 10 and Update equations 12 are applied repeatedly to solve for the system response over a period of time.

An additional type of mathematical model describes systems using algebraic equations as depicted in Figure 1C. The dynamic system 14 uses an algebraic equation 16 that needs to be solved at each time to obtain the outputs. While simple systems may allow one to obtain a closed-form solution for the system inputs and outputs, practical algebraic equations may best be solved iteratively using a numerical method involving both perturbations and iterations. Algebraic equation solving techniques used in the context of dynamic system modeling are discussed in greater detail below.

A fourth type of mathematical model is a composite system that has components that fall into the three types of models discussed above. Most complex real-world system models fall into this category. This class of systems has Output, Derivative, Update, and potentially other equations. Solving for the output response of such systems requires a combination of the solution approaches discussed for all of the classes above. One example of a composite system is one described by differential-algebraic equations (DAEs) which contain both differential equations and algebraic equations.

Grouped within the composite class of systems are many extensions involving relationships (equations) defined in terms of both outputs and state. For example, one can define a limited integration relationship for a differential variable. This relationship requires a set of equations that consists of the Output equation, an Update equation, a Derivative equation, and a Zero Crossing equation. The Zero Crossing equation defines the points in time where the upper and lower limits of the limited integration occur. Another example of an extension is the notion of Enable and Disable equations that define relationships among states or signals when parts of a system are activated and deactivated during execution.

Inherent in the four classes of systems (ODE, difference equations, algebraic equations and composite) is the notion of system sample time. The sample-time is the time interval at which the inputs, state, or outputs (collectively referred to as the results) of the system are traced as time progresses. Based on sample times, a system can be described as a discrete-time system, continuous-time system and hybrid system.

A discrete-time system is a system in which the evolution of the system results are tracked at finite intervals of time. In the limit as the interval approaches zero, the discrete-time system becomes a continuous-time system. The intervals of time may be periodic or non-periodic. Sometimes, non-periodic rate systems are referred to as non-uniform rate systems meaning that there is no periodic rate at which the response can be tracked. Nonuniform-rate systems can fall into the class of composite systems where an additional equation (`GetTimeOfNextVarHit`) defines when in the future the other equations associated with the system should be evaluated. A continuous-time system is a system in which the evolutions of the system results are continuously changing. Continuous-time signals change during numerical integration (minor time steps). An example of a continuous-time system is one described by an ODE. There can also be algebraic or composite continuous-time systems. A hybrid system is a system with both discrete-time and continuous-time elements.

If a system has only one sample time, it is said to be single-rate. If a system has multiple sample times, it is said to be multi-rate. Multi-rate systems can be evaluated (executed) using either a single-tasking form of execution or a multi-tasking form of execution. When multi-tasking execution is used, it conforms to rate monotonic scheduling principals as defined by Liu, C. L., and LAYLAND, J. W. *Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment*. ACM 20, 1 (January 1973), 46-61. Systems may also be categorized by the type of numerical integration solver being used. A fixed-step system is one that uses a fixed-step solver. Fixed-step solvers typically use explicit methods to compute the next continuous state at fixed periodic intervals of time. A variable-step system is one that is using a variable-step solver. A variable-step solver can use either implicit or explicit methods to compute the next continuous state at non-periodic intervals of time. Generally, variable-step solvers use a form of error control to adjust the interval size such that the desired error tolerances are achieved.

In practice, except for the most basic systems, mathematical models for dynamic systems involve a complex set of mathematical transformations applied in some prescribed manner with the outputs of some transformations forming the inputs of others. Each elemental transformation may be viewed in isolation as a simple dynamic system falling into one of the categories listed above. Therefore, a complex dynamic system may be modeled as an interconnection of various simple dynamic systems. A schematic representation of such an interconnection that has evolved over the years is the block diagram. Such block diagram models have now become a standard means in textbooks, design papers, journal articles, and specifications to communicate the details of a dynamic system's behavior.

A block diagram model of a dynamic system is represented schematically as a collection of blocks interconnected by lines that represent signals. A signal represents the input and output of a dynamic system. Each block represents an elemental dynamic system. A line emanating at one block and terminating at another signifies that the output of the first block is an input to the second block. Each distinct input or output on a block is referred to as a port. Signals correspond to the time-varying quantities represented by each line connection and are assumed to have values at each time instant. The source block of a signal writes to the signal at a given time instant when its system equations are solved. The destination blocks of this signal read from the signal when their system equations are being solved. The basic components of a block diagram are illustrated in Figure 2. The block diagram includes a plurality of blocks, lines and ports that are interconnected. Those skilled in the art will recognize that the term “blocks” does not refer exclusively to elemental dynamic systems but may also include other modeling elements that aid in readability and modularity of block diagrams.

The theory of Digital Signal Processing (DSP) focuses on modeling signals as sequences of samples. This view naturally fits into the time-based block diagram paradigm by mapping the samples $u[n]$ to discrete-time points $u(t_k)$. This adds the benefit of being able to model the interaction between DSP systems and other classes of time-based systems, e.g. continuous and/or discrete-time control systems.

Put another way, block diagram models are time-based relationships between signals and state variables representative of a dynamic system. The solution (computation of system response) of the model is obtained by evaluating these relationships over time, where time starts at a user-specified “start time” and ends at a

user-specified “stop time”. Each evaluation of these relationships is referred to as a time step. Signals represent quantities that change over time, and these quantities are defined for all points in time between the block diagram’s start and stop time. The relationships between signals and state variables are defined by sets of equations represented by blocks. These equations define a relationship between the input signals, output signals, state, and time. Inherent in the definition is the notion of parameters, which are the coefficients of the equations.

It is important to note that block diagrams are not exclusively used for representing time-based dynamic systems but also for other models of computation. For instance, flow-charts are block diagrams used to capture process flow and are not generally suitable for describing dynamic system behavior. Data flow block diagrams are block diagrams that describe a graphical programming paradigm where the availability of data (often thought of as tokens) is used to initiate the execution of blocks, where a block represents an operation and a line represents execution dependency describing the direction of data flowing between blocks. As used herein, the term block diagrams means time-based block diagrams used in the context of dynamic systems except as otherwise noted.

Block diagram modeling has spawned a variety of software products such as Simulink from The MathWorks, Inc. of Natick, Massachusetts, that cater to various aspects of dynamic system analysis and design. Such products allow users to perform various types of tasks including constructing system models through a user-interface that allows drafting block diagram models, allowing augmentation of a pre-defined set of blocks with custom user-specified blocks, the use of the block diagram model to compute and trace the temporal evolution of the dynamic system’s outputs (“executing” the block diagram), and automatically producing either deployable software systems or descriptions of hardware systems that mimic the behavior of either the entire model or portions of it (referred to herein as “code generation”). Each of the tasks listed above has many intricate details and subtle variations which are explored further below.

Block modeling software includes a number of generic components. Although the discussion contained herein focuses on Simulink version 5.0 (Release 13) from the MathWorks, Inc. of, Natick MA, those skilled in the art will recognize that it is applicable to other block modeling software applications. The generic components include a block diagram editor, blocks and a block diagram execution engine. The

block diagram editor allows users to perform such actions as draw, edit, annotate, save, and print out block diagram representations of dynamic systems. As noted earlier, blocks are the fundamental mathematical elements of a classic block diagram model. Simulink extends the classic block diagram models by introducing the notion of two classes of blocks, non-virtual blocks and virtual blocks. Non-virtual blocks are elementary dynamic systems. A virtual block is provided for graphical organizational convenience and plays no role in the definition of the system of equations described by the block diagram model. Examples of virtual blocks are the Bus Creator virtual block and Bus Selector virtual block which are used to reduce block diagram clutter by managing groups of signals as a "bundle". Virtual blocks may be used to improve the readability of models. Simulink further extends the meaning of a non-virtual block to include other semantics, such as a "merge" block semantic. The merge block semantic is such that on a given time step its output is equal to the last block to write to an input of the merge block. An additional extension provided by Simulink is the concept of conditional execution. Simulink contains the concept of conditional and iterative subsystems that control when in time block methods execute for a sub-section of the overall block diagram.

A block diagram execution engine contributes to the modeling software task of enabling the computation and tracing of a dynamic system's outputs from its block diagram model. An execution engine carries out the task of compiling and linking the block diagram to produce an "in-memory executable" version of the model that is used for generating code and/or simulating or linearizing a block diagram model. Note that execution of the block-diagram is also referred to as simulation. The compile stage involves checking the integrity and validity of the block interconnections in the block diagram. In this stage, the engine also sorts the blocks in the block diagram into hierarchical lists that are used when creating the block method execution lists. In the link stage, the execution engine uses the result of the compiled stage to allocate memory needed for the execution of the various components of the block diagram. The linking stage also produces block method execution lists which are used by the simulation or linearization of the block diagram. Included within the link stage is the initialization of the model which consists of evaluating "setup" methods (e.g. block start, initialize, enable, and constant output methods). The block method execution lists are generated because the simulation and/or linearization of a model must execute block methods by type (not by block) when they have a sample

hit.

After linking has been performed, the execution engine may generate code. In this stage, the execution engine may choose to translate the block diagram model (or portions of it) into either software modules or hardware descriptions (broadly termed code). If this stage is performed, then the stages that follow use the generated code during the execution of the block diagram. If this stage is skipped completely, then the execution engine uses an interpretive mode of execution for the block diagram. In some cases, the user may not proceed further with the execution of the block diagram because they would like to deploy the code outside the confines of the block diagram software. Upon reaching the simulation stage, the execution engine uses a simulation loop to execute block methods in a pre-defined ordering upon a sample hit to produce the system responses they change with time.

For linearization, Simulink uses the block method execution lists in a prescribed fashion to produce a linear state space representation of the dynamic system described by the block diagram.

The block diagram editor is the graphical user interface (GUI) component that allows drafting of block diagram models by a user. In Simulink, there is also a textual interface with a set of commands that allow interaction with the graphical editor. Using this textual interface, users may write special scripts that perform automatic editing operations on the block diagram. A user generally interacts with a set of *windows* that act as canvases for the model. There is generally more than one window for a model because models may be partitioned into multiple hierarchical levels through the use of subsystems (discussed further below).

A suite of GUI tools in Simulink allows users to draft a block diagram model on the corresponding windows. The GUI tools include a block palette, wiring line connection tool, annotation tool, formatting tool, attribute editing tool, save/load tool and publishing tool. The block palette is a *library* of all the pre-defined blocks available to the user when they are building the block diagram. Individual users may be able to customize this palette to: (a) reorganize blocks in some custom format, (b) delete blocks they do not use, and (c) add custom blocks they have designed. The palette allows blocks to be dragged through some human-machine interface (such as a mouse or keyboard) from the palette on to the window (i.e., model canvas). The graphical version of the block that is rendered on the canvas is called the icon for the

block. There may be different embodiments for the block palette including a tree-based browser view of all of the blocks.

The wiring line connection tool allows users to draw directed lines that connect the ports of blocks in the model's window. Lines are also added through various mechanisms involving human-machine interfaces such as the mouse or keyboard. Simulink also provides various forms of auto-connection tools that connect blocks automatically on user request to produce an aesthetically pleasing layout of the block diagram (especially those with high complexity with large numbers of blocks). The annotation tool allows users to add notes and annotations to various parts of the palette for a block diagram. The formatting tool enables users to perform various formatting operations that are generally available on any document editing tool. These operations help pick and modify the various graphical attributes of the block diagram (and constituent blocks) such as include font-selection, alignment & justification, color selection, etc. The block diagram and all the blocks within the block diagram generally have a set of functional attributes that are relevant for the execution or code-generation. The attribute editing tool provides GUIs that allows these attributes to be specified and edited. The save/load tool allows a created block diagram model to be saved. The saved model can be reopened in the editor at some later juncture through a load mechanism. Simulink also allow users to save blocks including pre-constructed subsystems into a separate class of block-diagrams called libraries. Such libraries facilitate reuse of the same block in a number of other block-diagrams. The load/save mechanism is specially equipped to handle loading and saving of blocks in a block-diagram that actually reside in libraries.

The publishing tool enables the viewing of the block diagram as a document that can be published in any of the standard document formats (examples: PostScript, PDF, HTML, etc.). Those skilled in the art will recognize that the windows for multiple models and all of the tools mentioned above could potentially be embedded in a single Multi-Document Interface (MDI) for providing a unified software environment. Those skilled in the art will also recognize that block-diagram packages offer scripting languages for writing out programs that automatically carry out a series of operations that would normally require interaction with the GUI. For example, Simulink offers a set of commands in MATLAB for carrying out operations such as block addition (*add_block*), block deletion (*delete_block*), starting and terminating execution (*set_param*), modifying block attributes (*set_param/get_param*), etc.

Simulink also offers a variety of other GUI tools that improve the ability of users to build and manage large block diagrams. Examples of such GUIs include: (a) a Finder that helps find various objects such as blocks and lines within a block-diagram, (b) a Debugger that helps debug the execution of block-diagrams, (c) a
 5 Revision Control UI for managing multiple revisions of the block-diagram, and (d) a Profiler for viewing timing results while executing a block-diagram.

· A typical base data-structure for a block may be represented as:

```

class Block {
10   public:
        // Access methods for setting/getting block data
        ...
        // Methods for block editing
        virtual ErrorStatus BlockDrawIcon();
15   virtual BlockParameterData BlockGetParameterData();
        ...
        // Methods for block compilation
        ...
        // Methods for block execution
20   .....
        virtual ErrorStatus BlockOutput()    = 0;
        virtual ErrorStatus BlockDerivative() = 0;
        virtual ErrorStatus BlockUpdate()    = 0;
        ...
25   private:
        BlockGraphicalData blkGraphicalAttributes;
        BlockFunctionalData blkFunctionalAttributes;
        BlockCompiledData   blkCompiledAttributes;
        BlockExecutionData  blkExecutionData;
30   ...
};
  
```

Although the example of the data structure above is written in C++, those skilled in the art will recognize that equivalent data structures written in other languages may also be used. The major data fields of the block data structure fall into four

categories, a graphical attributes field, a functional attributes field, a compiled attributes field and an execution data field.

The graphical attributes field is responsible for storing information relevant for graphical rendering of the block within its parent block diagram's GUI. Attributes specific to the block icon such as font, color, name, and icon-image are stored in this field. It should be noted that modifying these attributes does not affect the dynamics of the model using this block. The functional attributes field is responsible for specifying block attributes that may potentially affect the dynamics of the model using this block. These attributes are specified for the block as a whole and the input and output ports of the block. Examples of block attributes include block sample times and restrictive flags. Block sample times specify if the block corresponds to an elemental, continuous, discrete, or hybrid dynamic system. If the block is an elemental discrete-time system, then the attribute specifies the spacing between time instants at which the block response should be traced. A restrictive flag disallows the use of blocks in certain modeling contexts. For example, one may impose the restriction that there may only be one instance of given block in a model.

Attributes of block ports specify properties of the data that is either available or produced at that port. Block port attributes include dimensions, datatypes, sample rates, and direct feed through. Dimension attributes are individual dimensions of a multi-dimensional matrix that is used as a container for data elements. Datatype attributes are the datatype of each element of data in the data container. A complexity attribute is a flag to specify if each data element is real or complex. A sample rate attribute specifies how when the signal corresponding to an input or output port will be used. The port sample times may sometimes be used to implicitly infer the block's sample time. The direct feed through attribute is specified only for input ports and indicates whether or not the Output and/or GetTimeOfNextHit equations of the block are a function of the given input. This attribute helps in determining the sequence in which block methods should be executed while executing the block diagram.

The compiled attributes field of the block data structure holds the attributes of the block and its ports that mirror the functional attributes listed above. This field is filled in during block diagram compilation by utilizing the functional attributes of the block in conjunction with the functional and compiled attributes of the blocks that are connected to it. This process of determining the compiled attributes from the functional attributes is termed attribute propagation. Attribute propagation is

described in greater detail below, in the section on block diagram compilation. The execution data field is mainly responsible for storing the memory locations that are going to serve as sources for block inputs, outputs, states, parameters, and other work areas during execution of blocks.

5 The block data structure also has a set of associated methods that may be categorized as access methods to data fields, methods used in editing, methods used in compilation and methods used in execution. Access methods to data fields help in setting and getting the various data fields of the block. Methods used in editing are called by the block diagram editor in order to render the block appropriately in the
10 GUI of its parent block diagram. For instance, this set of methods may include a BlockDrawIcon method that determines the shape the block icon has on the GUI. Methods used in compilation are methods that are called by the block diagram compilation engine. They help validate the connections of the block to other blocks on the block diagram. The methods used in execution include a number of different run-
15 time methods that are required for execution. These include the BlockOutput, BlockUpdate, BlockDerivative methods that realize the Output, Update, and Derivative equations discussed earlier in the context of dynamic systems. In addition, to these methods Simulink includes several other run-time methods, such as the Jacobian, Projection, ZeroCrossings, Enable, Disable, Initialize, EvalParams (check
20 and process parameters), and GetTimeOfNextHit methods. It should be noted that there is no explicit method for algebraic equations because these are represented and processed in a different manner which will be discussed below.

 The base data structure for the block specifies the generic fields and interfaces that need to be supported by a block. Some of the methods are purely
25 virtual and have no specific implementation in the base block class. In order to define a specific block (such as an Integrator block), one needs to subclass the base block class and provide explicit definitions for these virtual methods. An example of the subclassing of a block may be seen by examining an Integrator block. Figure 3 depicts the desired behavior of an Integrator block 30. In order to create the subclass, four major categories of information within the subclass must be specified, the block
30 parameters, the methods used in editing, the methods used in compilation, and the methods used in execution.

 The elemental dynamic system embodied by the block may be parameterized as illustrated in Figures 1A-1C. Each block needs to be able to specify its list of

expected parameters. The block diagram editor's Attribute-Editing tool may allow users to specify the parameters for the block when they use it in their models. In the Integrator block example, the block has one parameter that specifies the block's initial condition for the block's state. Regarding the methods used in editing, the subclass

5 needs to specify a method that renders its icon. For example, the Integrator block may implement a method that makes its icon be a box with a '1/s' within the box. Also, the subclass needs to instantiate a method that allows access of the block parameters from the GUI's Attribute-Editing tool. For the Integrator example, this method would allow users to specify the Initial Condition parameter on a GUI for the block. For the

10 methods used in compilation, the subclass needs to instantiate methods that help in the compilation of the block diagram model in which it is placed. These methods help specify the compiled information for the inputs and outputs of the block. For instance, the Integrator block may specify a method that ensures that if the input to the Integrator is a vector, then the output is a vector of the same size. For methods used

15 in execution, the subclass needs to instantiate specific Output, Derivative, and Update methods that represent the block behavior. In the case of the Integrator block, an Output and Derivative method are needed. It should be noted that in Simulink the Integrator block has additional methods that are not illustrated here. The Output method sets the output to be equal to the state. The Derivative method sets the

20 derivative of the state to be equal to the input.

The specification of these four types of information for the Integrator block subclass may be shown by a reduced form of the Simulink Integrator block:

```

IntegratorBlock : public Block {
25   public:
       ErrorStatus BlockDrawIcon() {
           // Draw '1/s' on the icon
           .....
       }
30   BlockParameterData BlockGetParameterData() {
       // Return initial_condition as block data
       .....
   }

```

```

        ErrorStatus BlockOutput(){
            // Implement  $y(t) = x(t)$ 
            .....
        }
5      ErrorStatus BlockDerivative(){
            // Implement  $dx(t)/dt = u(t)$ 
            .....
        }
    private:
10     double initial_condition;

};

```

It should be noted that block diagram software generally provides open access to the
15 block's data structure to users of the software. This allows users to create and utilize custom block implementations in their models.

Blocks in a block diagram may be virtual or non-virtual. The designation of a block as non-virtual indicates that it influence the equations in the mathematical model for the dynamic system. In the context of block diagram software, it is
20 beneficial to include other *virtual* blocks that *do not* affect the equations in the dynamic system's model. Such blocks help improve the readability and modularity of the block diagram and wield no semantic influence on the mathematical model. Examples of such virtual blocks include virtual subsystems, inport blocks and output blocks, bus creator blocks and From and Goto blocks.

Modularity may be achieved in a block diagram by layering the block diagram through the use of subsystems. A subsystem facilitates layering by allowing a collection of blocks to be represented by a single block with input and output signals. The input and output signals of the subsystem are accessible to the constituent blocks within the subsystem. A subsystem is a virtual subsystem if its constituent blocks are
30 moved back into the main block diagram model during the model's execution. Within a virtual subsystem graphical entities, called inport and output blocks, are provided to define signal connections to the parent block diagram. These inport and output blocks indicate a tunnel-through signal connection to the parent block diagram.

Additional types of virtual blocks include bus creator blocks and selector blocks. In large models, there may be an extensive set of lines that connect one section of a block diagram to another section. To avoid excessive clutter of lines and improve readability, there is typically a special block called a *Bus Creator* that helps
 5 bundle all of the lines together to form a single *bus line*. This single bus line then connects the two sections of the model. At the destination end of the line, a block called a *Bus Selector* helps un-bundle the individual lines so that they can be connected to other blocks.

Other virtual blocks include From blocks and Goto blocks that are special
 10 blocks that help avoid graphical clutter, e.g. a line that connects two distant sections of a block diagram. The line is terminated close to its originating point by a From block. At the other end, a new line is drawn from a Goto block that is hot-linked to the From block. Each Goto and From block has an associated tag that describes which blocks are connected together. An important point to be noted is that Virtual blocks
 15 have neither execution data nor execution methods in their data structure.

Simulink also provides the user with the ability to extend the simulator by providing the ability to enhance the simulator with blocks that define dynamic systems or are virtual properties. The extension is provided through a language independent API (e.g. C, C++, Ada, Fortran, Assembly, M).

As noted previously, to facilitate modeling fairly large and complex dynamic
 20 systems, Simulink allows users to layer their block diagrams. A subsystem facilitates such layering by allowing a collection of blocks to be represented by a single block with input and output signals. The input and output signals of the subsystem are accessible to its constituent blocks. By nesting subsystems within each other, one can
 25 create block diagrams with arbitrary layers of hierarchy. Ideally a subsystem has no impact on the meaning of the block diagram. Additionally, subsystems provide a way of grouping blocks together and allowing other block diagram constructs to impose unified control on the constituent blocks. To enhance the modularity of subsystems, modeling software also allows aggregated list(s) of parameters of the blocks within
 30 the subsystem to be accessed from a single GUI, and defines and displays special icons on the subsystems. The process of defining the parameter list and the special icon is called masking a subsystem.

There are two main types of subsystem blocks, virtual subsystems and non-virtual subsystems. Virtual subsystems serve the purpose of providing the block

diagram with a graphical hierarchy. Non-virtual subsystems behave like an elemental dynamic system with its own execution methods (Output, Update, Derivatives, etc.). These execution methods in turn call the execution methods of the constituent blocks.

The classes of non-virtual subsystems are:

5

Atomic subsystems. These are similar to virtual subsystems, with the advantage of grouping functional aspects of models at a given layer. This is useful in modular design.

10 **Conditionally-executed subsystems.** These are non-virtual subsystems that execute only when a precondition is fulfilled:

Enabled subsystems. These are similar to Atomic subsystems, except that the constituent blocks only execute when an enable signal feeding the subsystem is greater than zero.

15 **Triggered subsystems.** These are similar to Atomic subsystems, except that the constituent blocks only execute when a rising and/or falling signal is seen on a triggering signal feeding the subsystem.

Enable with Trigger subsystems. These are an intersection of the properties of Enabled and Triggered subsystems.

20 **Action subsystems.** These subsystems are connected to action-initiator (e.g., an “If” or “SwitchCase” block), a block that explicitly commands the subsystem contents to execute. These subsystems are similar to Enabled subsystems except that the management of the “enabling” signal has been delegated to an action-initiator. Action subsystems define a new type of signal, called an action signal that signifies which subsystems are commanded to execute by the action-initiator.

25 **Function-call subsystems.** These subsystems provide a means of collecting blocks into a subsystem that is only executed when called by an owner block. The owner block may compute input signals for the subsystem before calling the subsystem. Additionally, the owner may also read output signals from the subsystem after calling it. Function-call subsystems define a new type of execution control signal, called a
30 function-call signal that contains no data. It is used to define the execution relationship between the owner block and the function-call subsystem. Function-call owners may also designate themselves as an “interrupt” source. In simulation, they simulate the effects of an interrupt and in code generation they can attach themselves to an (asynchronous) interrupt.

While subsystems and For subsystems. These subsystems execute the constituent blocks multiple times on a given time step.

Simulink allows for several forms of block parameters to be defined. There are two general categories of parameters: those parameters that can be modified during simulation and those that cannot be modified. An example of a parameter that may be modified during simulation is the amplitude of a Sine Wave block if configured by the user to allow modification during execution. A parameter such as the amplitude specifies coefficients of the dynamic equation, in this case the amplitude of the sine wave function defined by the Sine Wave block. An example of a parameter that can never be modified during simulation is the sample time of the Sine Wave block. The parameters that can be modified during simulation are further broken down into other categories which include mapping the dialog parameter (e.g. the amplitude) to run-time parameters or converting the dialog parameter to an inlined (non-modifiable) parameter. Run-time parameters can further be mapped to mathematical expressions of tunable Matlab variables or Matlab parameter objects describing properties of the variables (called Simulink.Parameter's). A global run-time parameter data structure is used within Simulink to manage the block parameters during the execution of the model.

In addition to block parameters, there are model-wide parameters which are generally associated with the solver. These parameters include aspects such as the time span in which to perform a simulation, the type of solver, and the time span. Simulink gives the user the ability to adjust solver parameters during model execution. The adjustment of these solver parameters is performed at the start of a time step.

Once a block diagram model has been constructed using the editor, an execution engine allows the model to be solved in order to trace the system outputs as a function of time. The solution of the model, which may be referred to as model execution, is carried out over a user-specified time span for a set of user-specified inputs. Simulation proceeds in four major stages: compilation, link, code generation, and the simulation loop. Alternatively, the execution engine can obtain a linear representation of the model (linearization). The interrelationship between the various stages is illustrated in a flowchart in Figure 4.

The execution begins when the block diagram 440 is compiled 442. Following the compilation stage, is the model link stage 444 which may also produce

linear models 446. Code may or may not be generated 445. If code is generated 448, a decision is made 449 whether to continue the simulation. If the decision is made to continue the simulation the model is simulated/executed through the Simulation Loop 450. If the simulation is not continued, the code may be delivered to a target 452 and
 5 executed in an external mode 454. If code is not generated the block diagram may execute in interpretive mode when entering the Simulation Loop 450.

The compile stage marks the start of model execution and involves preparing data structures and evaluating parameters, configuring and propagating block characteristics, determining block connectivity, and performing block reduction and
 10 block insertion. The preparation of data structures and the evaluation of parameters creates and initializes basic data-structures needed in the compile stage. For each of the blocks, a method forces the block to evaluate all of its parameters. This method is called for all blocks in the block diagram. If there are any unresolved parameters, execution errors are thrown at this point.

15 During the configuration and propagation of block and port/signal characteristics, the compiled attributes (such as dimensions, datatypes, complexity, or sample time) of each block (and/or ports) are setup on the basis of the corresponding functional attributes and the attributes of blocks (and/or ports) that are connected to the given block through lines. The attribute setup is performed through a process
 20 during which block functional attributes “ripple through” the block diagram from one block to the next following signal connectivity. This process (referred to herein as “propagation”), serves two purposes. In the case of a block that has explicitly specified its block (or its ports’) functional attributes, propagation helps ensure that the attributes of this block are compatible with the attributes of the blocks connected
 25 to it. If not, an error is issued. For instance, if an Integrator block is implemented to only accept numbers of double precision datatype, then this block will error out if it is driven by a block that produces single precision data, unless the user has asked for an implicit data conversion. Secondly, in many cases blocks are implemented to be compatible with a wide range of attributes. Such blocks adapt their behavior in
 30 accordance with the attributes of the blocks connected to them. This is akin to the concept of polymorphism in object-oriented programming languages. For instance, a discrete-time Filter block could be implemented to accept any of the standard integer datatypes ranging from 8-bit to 128-bit. The exact implementation of the block is chosen on the basis of the specific block diagram in which this block finds itself.

Included within this step are other aspects such as validating that all rate-transitions within the model yield deterministic results and that the appropriate rate transition blocks are being used.

5 The compilation step also determines actual block connectivity. Virtual blocks play no semantic role in the execution of a block diagram. In this step, the virtual blocks in the block diagram are optimized away (removed) and the remaining non-virtual blocks are reconnected to each other appropriately. This compiled version of the block diagram with actual block connections is used from this point forward in the execution process

10 Once actual block connectivity has been determined (by removing the virtual blocks) the block diagram may be further optimized by performing block reduction and insertion. During this step, non-virtual blocks may be inserted or a set of non-virtual blocks may be completely removed or reduced to a single equivalent block. Block insertion and reduction is mainly done to improve execution efficiency.

15 Examples of block insertion and reduction include the removal of Gain blocks whose gain value is 1. A Gain block is a block which multiplies its input value by a gain parameter, such as a simple amplifier. Figure 5 depicts the replacement of a collection of blocks 460, 462, and 464 connected in an accumulator pattern and leading to result 66 with an equivalent synthesized block 468 representing the

20 accumulator pattern leading to the same result 466. A signal copy block may also be automatically inserted in order to make contiguous memory copies of signals that are made up of disjoint memory sections. Block insertion and reduction may also be performed at other suitable stages of compilation.

25 The way in which blocks are interconnected in the block diagram does not necessarily define the order in which the equations (methods) corresponding to the individual blocks will be solved (executed). The actual order is partially determined during the sorting step in compilation. Once the compilation step has completed, the sorted order cannot be changed for the entire duration of the block diagram's execution.

30 The first step in sorting involves transforming the graphical block diagram into a compiled (in-memory) directed graph consisting of arcs and vertices. The vertices are derived from some of the non-virtual blocks. For instance, virtual and reduced blocks do not appear in the directed graph. The arcs represent data dependencies between the vertices. The data dependencies do not correspond to the

signals in the block diagram. For example, all signals that connect to input ports without direct feed through are “cut” or ignored. In addition, data dependencies are added to capture implicit dependencies. For example, all inputs to a Function-Call subsystem are implicit data dependencies to the owner (caller) block.

5 The process of converting a block diagram into a compiled directed graph is shown in Figure 6A. A block diagram 481 includes a Sine Wave 1 block 482, a Sine Wave 2 block 484, a Goto block 486, a Function Call Generator block 488, and a From block 490. Also included are a Function Call Subsystem block 492, a Sum block 494, a Gain block 496, an Integrator block 498 and an Output (Output 1) block
10 100. Those blocks which are not virtual or reduced appear on the corresponding directed graph 111. The directed graph 111 includes a Sine Wave1 vertice 112, a Sine Wave 2 vertice 114, a function-call generator vertice 116, and a function call subsystem vertice 118. Also included are a Sum vertice 120, a Gain vertice 122, an Integrator vertice 124 and an Output 1 vertice 126. The vertices are connected by
15 arcs.

 The graph is used to sort the blocks into a linear sorted list. Figure 6B depicts a sorted list 128 generated from the compiled directed graph 111 which includes the elements appearing as vertices in the directed graph 111 sorted into order. The root block diagram has a sorted-list associated with it. Roughly speaking, each non-virtual
20 subsystem layer and some special block diagram elements also each have their own sorted-list. During the sorting of the graph into the list, strongly connected components are identified. The term strongly connected section, which is a term that originates from graph theory, is a subset, S , of the blocks of a block diagram such that any block in S is reachable from any other block in S by following signal connections
25 and S is not a subset of any larger such set. Strongly connected sections are flagged as algebraic loops when all blocks have direct feed through (an example is shown in Figure 6A consisting of the Sum 120 and Gain 122 blocks). Such loops correspond to a set of algebraic equations and are solved using iterations and perturbations during block diagram execution by solving for the algebraic variables. Algebraic variables
30 are either specified by the user via Initial Condition blocks or chosen by the execution engine. Solving of algebraic loops is discussed further below.

 Sorting must also take into consideration other user specified dependencies between the blocks. These dependencies include the concepts of priorities and placement groups. A block priority specifies the order in which the equations

associated with a block are evaluated with respect to other blocks. Placement groups are a way of causing each class of block methods for a specified set of blocks to be “placed together” in the block method execution lists. The terms “data dependency” or “data precedence” as used herein refers to the arcs of the compiled directed graph and not the signals found within a block diagram. Attempting to correlate data dependencies directly to the signals found within a block diagram is incorrect and leads to the conclusion that Simulink does not satisfy data dependencies, i.e., the execution of the operations or block methods does not satisfy data dependencies if one interprets signal connectivity as specifying data dependencies.

After compilation, the link stage commences. During this stage physical memory allocations are made in order to prepare for execution. Buffers are allocated for block input and output data buffers, states, and work areas. Additionally, block method execution lists that are derived from the sorted list allow for execution of the block diagram. Each block method execution list is a list of block methods that are to be executed in a sequence when each method within the list has a sample hit. There is generally a set of block method execution lists associated with each layer of the block diagram that corresponds to a non-virtual subsystem. Non-virtual subsystems are either defined by the user or automatically synthesized during compilation to either efficiently execute the model or simplify the implementation of the semantics defined by Simulink. In multi-tasking mode, the lists within each layer may be further partitioned when block diagrams have blocks with different sample rates. These lists are explained in greater detail below.

Those skilled in the art will recognize that while the block method execution lists are derived from the sorted list, they do not necessarily correspond one-to-one with the sorted lists. First, each block method execution lists contains only blocks that have such a block method of the given type (class) defined by the list. Second, block methods corresponding to components like the function-call subsystem do not appear on the block method execution lists because they are executed by an “owner” block.

Although included in the discussion of the compilation stage, it is not required that the time-based diagram perform the block sorting step during compilation. The sorting step is performed to achieve efficient execution. Ignoring efficiency, there is no semantic reason to perform the sorting step. Any random ordering of the block methods will work. In fact, any ordering of all block method execution lists except the Output block method execution list will result in the same level of efficiency.

Randomly re-ordering the Output block method execution list will yield correct answers. If the Output block method list is randomly ordered, then the Simulation engine, when executing the Output block method execution list, continues sequencing through the Output block method execution list at each point in time until there are no changes.

Similarly included within the linking stage for the sake of simplicity, is the memory initialization of the model. The memory initialization of the model includes invoking block start, initialize, constant initialize, enable, and constant output methods. These are examples of some of the block methods that are used during model setup (prior to execution) to initialize the “state” of the system so that execution or linearization can commence.

The compiled and linked version of the block diagram may be directly utilized to execute the model over the desired time-span. This interpretive mode of execution is suitable for getting fine-grained signal traceability. It should be noted that the traceability associated with interpretive execution comes at the price of increased overhead in the form of additional execution-related data-structures and messaging in the engine. An alternative to the interpretive execution mode is to utilize the generated-code created by Real-Time Workshop tool for Simulink models. In this mode, the engine (upon the behest of the user) translates a selected portion of the block diagram (or the entire block diagram itself) into code. Such code could be in a number of possible forms. The code may be instructions in a high-level software language such as C, C++, Ada, etc., hardware descriptions of the block diagram portions in a language such as HDL, or custom code formats suitable for interpretation in some third-party software. Alternatively, the code may be instructions suitable for a hardware platform such as a microprocessor, microcontroller, or digital signal processor, etc., a platform independent assembly that can be re-targeted to other environments, or just-in-time code (instructions) that corresponds to sections of the block diagram for accelerated performance.

The execution of a portion of the block diagram represented in code may be performed in a number of different ways based on the specific code format. The portion of the block diagram may execute a compiled version of the code generated in a high-level language (accelerated or software-in-the-loop simulation), the execution may simulate code that corresponds to a hardware description on a hardware simulator, (co-simulation execution), the execution may involve calling out to third-

party software to run code generated for such software (co-simulation execution), or the execution may call out directly to hardware that will run code that was generated and compiled for that hardware (processor-in-the-loop execution).

There are several different advantages to execution through code generation:

- 5 Execution of generated code can be more efficient than interpretive execution because of fewer data-structures and lesser internal messaging in the engine, although the increased efficiency generally comes at the cost of decreased execution traceability. Simulation of hardware descriptions during execution can help identify and resolve bugs in the software stage of a design project. Such bugs prove much more expensive to track and fix once the system has been implemented in hardware. Additionally,
- 10 block diagram modeling software can be integrated with other software environments that are suitable for modeling and simulating special classes of systems. Models can be tested directly in hardware thereby making prototyping of new systems fast and cost-effective. For instance, consider the design of a controller for an anti-lock
- 15 braking system of a car. The dynamics of the braking system can be executed in the interpretive mode in the block diagram. The controller itself can be implemented on a hardware micro-controller to test the efficiency of the control laws implemented within. Note that for such target execution, it is normally necessary for the time span over which a model is executed by the software to match real-world time. In other
- 20 words, the software must allow real-time execution of the block diagram model. Those skilled in the art will recognize that when users generate code, they may choose to not proceed further with the block diagram's execution. They may choose to take the code and deploy it outside of the confines of the modeling software environment. This is normally the last step in the design of dynamic systems in a block diagram
- 25 software package.

- There are several forms of target code execution known to those skilled in the art such as Rapid Prototyping, Embedded System Deployment, and Hardware-in-the-Loop which execute a model or portions of a model via the generated code on a Real-Time System target. One aspect of deploying (executing) the generated code on a
- 30 target is the notion of "external mode." External mode refers to a system where Simulink acts as a monitor and debugger of the generated code running in real-time on a target. In External Mode, users can change parameters and view signals via standard Simulink elements. Another important aspect of the code generation technology is that it is very extensible. Provided with the Simulink product family is

the Target Language Compiler (TLC). This technology enables the creation of “active scripts” that control how the generated code is produced for a block diagram. Using TLC, one can tailor the generated code to suite their specific needs.

The execution of the block diagram uses a Simulation Loop (SimLoop) for solving for the block diagram’s outputs for a specified set of inputs over a specified span of time (“Time” in reference to the Simulation Loop means the time-line corresponding to the tracing of the dynamic system’s outputs, not real-world time unless otherwise noted). The term “SimLoop” applies to real-time systems where each iteration is tied to a physical periodic clock or other timer source. During this process, the block methods (equations) corresponding to the individual blocks are executed by type following their sorted order when they have a sample hit. The term “block execution” is loosely used to mean executing all block methods associated with the given block for a given time step, generally starting with the output method. Strictly speaking, blocks do not execute; the engine executes (evaluates) the appropriate block methods at the appropriate time points.

SimLoop has two variants “single-tasking” and “multi-tasking” depending on sample times. In general, the sample time of a block is the interval of time between calls to the Output, Update, and/or Derivative methods for a given block. In computing this interval, repeated calls at the same time instant (not in real-world time but the time corresponding to the execution of the dynamic system) are counted as the same call. A block’s sample rate may also be thought of the interval between successive executions of the block methods. If there is no uniform or regular interval between calls, then the block is said have a continuous sample time. If a uniform time interval can be found, then the block is said to have a discrete sample- time equal to that interval. Although blocks may be associated with more than one sample time in a sufficiently complex dynamic system the descriptions contained herein are confined to blocks with a single sample-time. Those skilled in the art will recognize that the descriptions may be extended to encompass blocks with multiple sample times.

Figure 7A depicts an abstract example of a block diagram being executed. The diagram includes a plurality of blocks 140, 142, 144, 146, 148 and 150. The block ports that have direct feed through are explicitly marked (using the symbol ‘df’) 152. Additionally, an abstract view of the execution methods instantiated by each block is shown in Figure 7B. The blocks contain a number of different methods 160, 162, 164, 166 and 168. Execution methods include the three basic execution methods

discussed earlier: Output, Update, Derivative, as well as several other methods that aid in advanced block functions such as initialization, linearization and zero-crossing detection (which are discussed below). The data-dependencies between the compiled vertices created during sorting are used to generate the Sorted List 170 shown in

5 Figure 7C.

A block diagram consisting of blocks that all have the same sample time is said to correspond to a single-rate system. A block diagram consisting of blocks that have more than one sample time corresponds to a multi-rate system. Figure 8 depicts a multi-rate system, adding sample-time information to the block diagram of Figure
10 7A. The plurality of blocks 140, 142, 144, 146, 148, and 150 each have an associated sample time. Since the sample times in the block diagram differ between blocks, the system is considered a multi-rate system. Block A 140, block E 148 and block F 150 each have a sample time of 0.1 seconds. Block B 142, block C 144 and block D 146 each have a sample time of 1.0 seconds.

15 The SimLoop is the heart of the execution engine. Each full pass through the loop is responsible for computing the outputs of the system at a particular time. At the end of each loop, the execution time corresponding to the next pass through the loop is computed. If this time exceeds the stop time specified by the user, the execution terminates. Within the loop, the sequence in which individual block equations are
20 solved is determined by two pieces of information: the sample times of the blocks and the sorted order determined during the Compile stage. The amalgamation of these two pieces of information gives the execution lists for the block diagram. Those skilled in the art will recognize that the execution lists are created in the Link stage and are explained in the context of SimLoops for convenience. There are two distinct
25 approaches for building execution lists and using them in the SimLoop. These approaches correspond to the Single-tasking and Multi-tasking SimLoops summarized in the discussion on Figure 10 below.

Simulink also has the ability to modify coefficients (parameters) of blocks who declare their parameters as tunable. An example of a block is a Sine Wave block
30 that implements the function $\text{output (time)} = \text{Amplitude} * \sin(\text{frequency} * \text{time} + \text{phase}) + \text{bias}$, where time is the independent variable and the parameters are: amplitude, frequency, phase, bias. When these parameters are declared as tunable, Simulink lets the user change these coefficients during simulation. Changing parameters is a drastic operation in that the definition of the model has changed (e.g.

the sine block defines equations that describe the system). Thus, to enable the changing of parameters during the SimLoop, Simulink first queues parameter changes and then applies them on the next time step. Thus, the changing of parameters is not immediate. The delay in the changing of parameters is needed to ensure system stability. The application of the parameters at the start of the next time step is combined with the reset of the solver (Integrator) if needed.

For the purpose of exploring single-task loops and multi-task loops, Figure 9 depicts the block diagrams of Figure 7A and Figure 8 where Method1 corresponds to the Output method 190 and Method2 corresponds to the Update method 192. All other methods are ignored in the explanation of the loops. Simpler loops which do not include blocks that have continuous sample times are used in the example since the explanation is simpler in the context of discrete sample times and it is straightforward to extend to continuous sample times.

In a single-tasking SimLoop, there is essentially a single execution time-line. On this time-line, each block is executed when it has a sample hit. A sample hit is defined to an execution time instant that is an integer multiple of the block's sample time. To aid in execution, execution lists are constructed for each method type. Figure 10 depicts the sequence of steps followed by a single-tasking execution loop. Following initialization (step 200), a time parameter is checked to see if the current time is less than the stop time (step 201). If the time is not less than the stop time, the simulation ends (step 202). If the time is less than the stop time, the simulation continues and the root output method execution list is executed (step 204). Following execution of the output method list (step 204) the update method execution list is executed (step 206). Following the performance of an integrate step (208) (the Integrate step is described more below in Figure 14), the time parameter is incremented by the applicable step size (step 210).

Blocks are arranged in the single-tasking execution lists in the sorted order as shown in Figure 11A. A sorted list 250 is used to generate an Output method execution list 252 and an Update method execution list 254. Referring back to the example in Figures 7 and 8, the engine sequentially steps through and execute each block in the block method execution list when the execution time divided by the sample time equals an integer number (1, 2, 3, 4, etc.). At time zero (T_0), all the blocks are executed. This involves executing the Output methods for blocks F, E, D, A, B, and C (in this order as dictated by the sorted list) and then executing the Update

methods of blocks F, E, and D (again, in this order based on the sorted list). The execution time then is then incremented by step size, which in this case is assumed to be 0.1 seconds. Execution then commences once again at the top of the loop for $T=0.1$ ($T_{0.1}$). Blocks F and E have a sample time of 0.1 seconds and have a sample hit
 5 ($0.1 \div 0.1 = 1$, sample time is an integer multiple of the execution time), so the output block methods for Blocks F and E are executed. Block D, however, has a 1.0 second sample time and has no sample hit ($0.1 \div 1.0 = 0.1$, sample time is not an integer multiple of the execution time), so its output block method is not executed (essentially it is skipped). Block A, like Blocks F and E, has a 0.1 second sample time and so its
 10 output block method is executed. Blocks B and C, like Block D, have 1.0 second sample times and are skipped during this iteration of the simulation loop, which completes execution of the output block method execution list for $T_{0.1}$.

The execution timing of the example block diagram in single task mode is shown in the first time-line of Figure 11B. In this diagram, note that the execution-
 15 time is not synchronized with real-world time. Instead, execution time progresses as fast as it can in real-world time. The sorted list 259 is executed on the time-line 260. The methods in the list 262 are executed at the appropriate time step 264. Block diagram modeling software can also allow users to simulate real-world conditions by synchronizing execution time with real-world time. Such execution is illustrated in the
 20 second timing diagram of Figure 11B. The methods 262 are implemented at a time-step 264 synchronized with real world time on the time line 270.

In multitask mode, the engine performs execution along multiple time-lines based upon the number of block sample times used in the mode as shown in the flowchart of Figure 13. In the example of Figures 7 and 8, the model's blocks have a
 25 sample time of either 0.1 seconds or 1.0 second. This implies that the engine runs one set of blocks along a 0.1 second time line and another set of blocks along a 1.0 second time line. In order to run in multitask mode, the execution lists are first divided on the basis of methods (as in single-tasking mode) and then subdivided again based upon block sample times. This is illustrated in Figure 12A. The sorted list 280 is used to
 30 generate an output method execution list 282 and update method execution list 288. The output method execution list 282 is split into two separate list execution lists 284 and 286 based on sample times. Similarly, the update method execution list 288 is divided into two update method execution lists 290 and 292 based on sample times.

The execution engine uses the divided execution lists to create multiple execution time lines. In the multitask mode the engine places a higher execution priority on the faster sample time blocks than the slower sample time blocks. This prioritization is carried out by assigning Task Identification Numbers (TIDs) to each execution list; the higher the priority, the lower the TID. For example, a TID of 0 executes at a higher priority than a TID of 1, and so forth. Furthermore, because, during execution in multitask mode, execution transitions between the faster and slower blocks, and vice- versa, the multitask mode requires rate transition blocks that allow the model to transition from blocks running at fast sample times, in our example 0.1 seconds, to slower samples times, e.g., 1.0 seconds. The rate transition blocks are required to correctly simulate how a multi-rate system would behave in a real-time environment. To provide this transition, the engine promotes rate transition blocks to the TID of the fast block for which transition is being provided, although the engine executes these rate transition blocks at their slower rate. This is why Blocks D and B appear in the 0.1 sample time output method execution list in Figure 12A.

The execution of our example in the multi-task mode may be seen in Figure 12B. At time $T=0$, the engine first executes the high priority output methods (those with TID 0) for Blocks F, E, D, A and B, then it executes the high priority update methods (those with TID 0) for Blocks F and E. After finishing the high priority blocks, the engine executes the lower priority output block methods (those with TID 1) for Block C, and then executes the lower priority update methods (those with TID 1), which, in this example, is Block D. In contrast to the single task mode, in multitask mode the engine runs through a TID inner loop to execute the output and update block methods before going on to the Integration step, as the flow chart in Figure 13 which is discussed below illustrates.

As a result of the inner TID loop, as well as the segregated block method execution lists, the order of execution in multitask mode differs from the order of execution in single task mode. Recall for the example that in single task mode that the order of execution at $T=0$ is: F_o , E_o , D_o , A_o , B_o , C_o , F_u , E_u , and D_u , where the subscript "o" stands for output method and the subscript "u" stands for update method. In the multitask mode, however, the order of execution at $T=0$ is: F_o , E_o , D_o , A_o , B_o , F_u , E_u , C_o , and D_u . Notice that C_o is executed in a different order in multitasking mode. This occurs because separate method execution lists (based upon sample time) are created and run in order from fastest sample time to slowest sample time.

Additionally, the use of rate transition blocks restricts the connection of blocks with different rates. By requiring the insertion of these blocks into the model, the engine ensures that execution in multitask mode will follow the sorted list.

After it is finished executing the block methods for $T=0$, like in the single task mode, the execution time is incremented (again assume by 0.1 seconds) and execution goes to the beginning of the loop. The engine executes F_o , E_o , A_o , F_u , and E_u , and the engine does not execute the block methods of Blocks D, B, and C because the current execution time is not an integer multiple of those block's sample time. The engine repeats this execution until the execution time is incremented to 1.0 seconds, whereupon execution occurs in the same manner as described for $T=0$. The engine repeats this overall process until the execution stop time is reached.

Figure 12B shows two time-lines; the lower time-line 306 represents the execution order of the faster sample time blocks (Blocks A, E, and F), along with the rate transition blocks (Blocks B and D), while the top time-line 308 shows the execution order of the slower sample time block (Block C), and the rate transition (Block D) update method. The time-lines are generated from the sorted list 302 and the associated sample times 304. The lower line, representing the faster sample times has a TID of 0, and the top line has a TID of 1. For execution time $T=0$, the chart shows that the engine executes the output methods for Blocks F, E, D, A, and B (designated on the chart as F_o , E_o , D_o , A_o , B_o). Then, consistent with the flow chart for the multi-tasking mode (see Figure 13 discussed below), the engine executes the update block methods for Blocks F and E (designated F_u , and E_u). Once the engine is finished with the high priority block methods, the output method for Block C (C_o) and the update method for rate transition block D (D_u) are executed. The execution time is then incremented by the step size (continue to assume 0.1 seconds) and the blocks that have a sample hit are executed. The figure shows execution of F_o , E_o , A_o , F_u , and E_u , which is repeated, as noted above, until execution time equals 1.0 second. Notice, like in the non-real-time case for Single-task mode, the engine does not wait for time to elapse; rather it executes block methods immediately upon completion of the previous pass through the loop.

Figure 13 shows the overall sequence of steps taken by Simulink in multitask mode. Following initialization (step 220), the output method execution list is executed for the fastest sample time (step 222). The update method execution list is then executed for the fastest sample time (step 224). A time parameter is checked

(step 225) to determine to determine if the time is less than a designated stop time. If the stop time has been reached, the simulation completes (step 226). Otherwise, the integrate stage (step 228) is performed. The task ID variable is incremented (step 230) and compared to a parameter of the number of sample times (step 231). If the task ID is less than the number of sample times, the output method execution list for the methods assigned the new task Id are executed (232) followed by the execution of the update method execution list assigned the new task ID (step 234). The task ID variable is incremented and the process iterates with the task ID being compared to the number of sample rate times (step 231). When the task ID number is determined to equal the number of sample rate times, the simulation time is incremented (step 238) and the entire process iterates with the output method list execution list (step 222) being executed for the fastest sample times. The process continues until the end of simulation when the time equals the stop time (Step 226).

15 In order to understand how the step size is picked within SimLoop, it is first necessary to understand the notion of a solver. The solver is a module of the execution engine that is responsible for performing two tasks: (a) determining how far execution time should be advanced between consecutive passes through the SimLoop in order to accurately trace the system's outputs, and (b) integrating the derivative of the states of the system to obtain the actual states. Based on how solvers perform the first task, they are generally classified into two basic classes: Fixed-step solvers or Variable-step solvers.

25 Fixed-step solvers are solvers in which the time step-size between consecutive passes through the SimLoop is a fixed quantity. The user generally explicitly specifies this quantity. These solvers are used to model types of systems that must operate within a defined time (discrete systems). For instance, an anti-lock braking system may be designed to control a car's braking system, and to execute such control in one-hundredth (0.01) of a second so as to assure the car stops safely; if the braking system does not meet its timing constraints, the car may crash. Fixed-step solvers, therefore, are designed to help model discrete systems that have to generate a result in a fixed time period, and the fixed-step execution assures that the modeled system can generate such results.

Variable-step solvers are designed to model continuous systems where non-evenly spaced time steps are needed to simulate all significant behavior. For

example, one may want to simulate the path of a bouncing ball, where it bounces, how high it bounces, and where it stops. It is known, based on experience, that the ball's bounces will not be evenly spaced, and that the height of the bounces will diminish as a result of gravity, friction, and other forces. Variable-step solvers are used for these types of continuous systems and to determine what step size to use so that the behavior of the ball will be accurately modeled.

The two broad classes of solvers are further subdivided based on the integration task they perform. There are several algorithms for carrying out numerical integration. The particular choice of the integration algorithm gives rise to the subclasses of solvers.

The difference in the conceptual definition of Fixed- and Variable-step solvers leads to the functional difference in the context of the SimLoop. The major difference between the solvers arises in the Integrate step of the SimLoop which is depicted in Figure 14. During the Integrate step, the Variable-step solver executes the Output and Derivative block method lists for a number of iterations that varies based on the solver subclass (i.e., the numerical integration algorithm it uses) and integration error tolerances. In a fixed-step solver, the number of iterations is fixed for a given solver subclass. Another difference between solvers arises in the Integrate phase in the context of an operation known as zero-crossing detection. Zero-crossings in the derivatives of the state generally indicate a discontinuity in the states themselves. Because discontinuities often indicate a significant change in a dynamic system, it is important to trace the system outputs precisely at such points. Otherwise, the outputs of the model could lead to false conclusions about the behavior of the system under investigation. Consider, again the example of the bouncing ball. If the point at which the ball hits the floor occurs between simulation steps, the simulated ball appears to reverse position in midair. This might lead an investigator to false conclusions about the physics of the bouncing ball. To avoid such misleading conclusions, it is important that the execution has time steps on and around the vicinity of discontinuities.

In the case of Fixed-step solvers, there is no notion of zero-crossing detection and one is not guaranteed to find all points of discontinuity. One can only keep reducing the step-size to increase the probability of hitting the discontinuity. Contrastingly, in the case of Variable-step solvers, the Integrate step explicitly includes zero-crossing detection. The execution step size is then adjusted accordingly

to ensure that discontinuities are tracked accurately. To enable zero-crossing detection, blocks that can produce discontinuities instantiate a special execution method. This method registers a set of zero-crossing variables with the execution engine, each of which is a function of a state variable that can have a discontinuity.

- 5 The zero-crossing function passes through zero from a positive or negative value when the corresponding discontinuity occurs. During the zero-crossing detection phase of the Integration step, the engine asks each block that has registered zero-crossing variables to update the variables for the projected time of the next time step. These variables are then checked for a change of sign since the current step. Such a
- 10 change indicates the presence of a discontinuity. An iterative process then tries to narrow down the location of the discontinuity and ensure that the next few time steps (at least 2) accurately bracket the location of the discontinuity. The final difference, which is in the step-size during execution, is a direct consequence of the two previous differences in the step-size determination. In Fixed-step solvers, the step size is a
- 15 known and fixed quantity. For Variable-step solvers, the step size is determined during the integration iterations and the zero-crossing detection that happens during the Integration step.

- An example of the variable-step solver is shown in Figure 14, the derivative method execution list is executed (step 240) followed by the output method execution
- 20 list (step 242). The derivative method execution list is then executed again (step 244) and the solver iterates between the execution of the output method execution list (step 242) and the execution of the derivative method execution list (step 244). A similar iteration loop then occurs between the execution of the output method execution list (step 246) and the execution of the zero-crossing method execution list (step 248).
 - 25 Note that Simulink also includes other methods such as Projections and Jacobians in this step as needed.

- While it is theoretically possible to have Variable-step solvers in the context of multitasking, such a combination is not employed in practice. This is because the step-size for such solvers can become very small making it impossible to keep up with
- 30 the real-time constraint that generally goes along with multitasking execution. An added complication is that the integration step in such solvers is iterative and takes varying amounts of time at each step of the execution. Therefore, Variable-step solvers are generally used only in conjunction with the Single-Tasking SimLoop.

Additionally, they are not usually employed in systems that need to operate in real-time.

When a model contains an algebraic loop, the engine calls a loop solving routine at each time step. The loop solver performs iterations and perturbations to determine the solution to the algebraic condition (if it can). One possible approach to solving the algebraic equation $F(z) = 0$, is to use Newton's method with weak line search and rank-one updates to a Jacobian matrix of partial derivatives. Although the method is robust, it is possible to create loops for which the loop solver will not converge without a good initial guess for the algebraic states z . Special blocks are generally provided to specify an initial guess of the states in the algebraic loop.

In addition to the various forms of the SimLoop, modeling packages such as Simulink use the output of the Link stage to compute linear models through a process generally referred to as model linearization. These linear models may be used in the SimLoop at various points in the execution of the overall model. Alternatively, the linear model may be returned to the user. The linearization process involves the use of a Jacobian method defined on blocks and numerical Jacobian algorithm.

Information related to the compiled block diagram may be presented to users in an automatically generated report. This report allows users to quickly obtain documentation of the functional description of their model. Information related to the execution of a particular model (such as the time taken to execute various portions of the model and the coverage of various portions of the model) may be obtained automatically and presented to the user as a report.

In a block diagram environment if all of a block's inputs, states and outputs change continuously or at one rate, the selected block is considered single rate. For example, all of the inputs, states and outputs of a block change at a rate of 5Hz. If the inputs, states, and outputs of a block change together or separately at two or more time intervals, such a block is considered multi-rate. Should a block diagram model include multi-rate blocks or two or more single-rate blocks operating at different rates, for example, 5 Hz and 10Hz, then the block diagram model is considered multi-rate.

Some block diagram environments may include an automatic code generator, such as an execution engine that translates a block diagram to executable entities following the layout of the block diagram. The executable entities are compiled and executed on an electronic device such as a computer to implement the functionality

specified by the model. Typically, the code generation preserves a model hierarchy in a call graph of the generated code. For instance, each subsystem of a model in a block diagram environment can map to a user specified function and the generated code.

5 One code generation scheme in a block diagram environment is referred to as “rate guarding”. A rate guarding code generation scheme generates a single top level function for the entire block diagram model. The single top level function operates as an entry point for each block in the model regardless of an operating rate of the block. The single top level function is called by each operating rate in the block diagram
10 model. That is, the identity (ID) of the calling rate is passed as an argument to the top level function. The top level function accepts the ID of the calling rate as an argument and passes the argument to lower level functions that correspond to subsystems and the block diagram model. Accordingly, code inside the functions, for example, one or more conditional statements, are conditionally executed depending
15 on the ID of the calling rate as passed as an argument to the function. That is, within each function, the block of the block diagram model corresponding to the ID of the calling rate is executed conditionally against the calling ID.

Often, the complexity of block diagram models results in a decreased rate of execution as the blocks in a block diagram model pass arguments identifying
20 operating rates and functions conditionally execute based on the passed operating rate ID. Accordingly, execution of the block diagram models based on passing an argument identifying an operating rate of an operation in the model becomes burdensome particularly in instances where the block diagram model is meant to
simulate a real-time system.

25

Summary of the Invention

The illustrative embodiment of the present invention enables implicit identification of an operating rate by a function in a block diagram environment. The illustrative embodiment provides an approach that avoids the need to pass an
30 argument to a function in order to identify an operating rate of an operation in a block diagram environment. The illustrative embodiment allows the generation of a top level function for each operating rate used in a block diagram model. Each top level function contains or is implicitly associated with blocks or components within blocks executing at the corresponding operating rate. The illustrative embodiment further

enables sub-systems within a block diagram model to request generation of separate functions, i.e. sub-functions for each operating rate used by the blocks of that sub-system.

5 In one aspect of the present invention a method for generating code from a multi-rate block diagram model in a block diagram environment is performed in an electronic device providing the block diagram environment. The method includes steps of identifying from the multi-rate block diagram model a number of operating rates contained therein. Each block in the multi-rate block diagram model has one or more operating rates and each of the identified operating rates indicate a rate of an associated operation of one of the blocks. The method includes a step of generating
10 code from the multi-rate block diagram model. The generated code having one function for each identified operating rate in the multi-rate block diagram model. Each of the functions provides implicit identification of one of the identified operating rates per groups of blocks in the multi-rate block diagram model. The method further
15 includes a step of generating code from the multi-rate block diagram model for sub-systems in the multi-rate block diagram model. The sub-systems include one or more element blocks having tasks operating at two or more rates. Generation of code for the sub-systems generates a number of sub-functions that correspond to one of the identified operating rates of the element blocks in the sub-systems in the multi-rate
20 block diagram model. The sub-function provides implicit identification of the operating rate for a sub-system of the groups of blocks in the multi-rate block diagram model. Each of the functions and sub-functions generated from the multi-rate block diagram model invokes only the blocks from the multi-rate block diagram model executing at the rate to which the function corresponds.

25 In another aspect of the present invention, a method is performed in an electronic device that provides a block diagram environment. According to the method a first function is instructed to invoke a portion of a first block from the block diagram executing at a first rate. A second function is instructed to invoke a portion of a second block from the block diagram executing at a second rate. The first and
30 second functions unconditionally define an execution path for each of the rates. The method is further capable of instructing a third function to invoke another portion of the first block from the block diagram executing at the first rate. The third function relates to the first function where the third function implicitly corresponds to a sub-

system of the first block. In this manner, the functions are associated with an argument list having selected arguments necessary for invoking the functions.

In a further aspect of the present invention, a method is practiced in an electronic device that provides a block diagram environment. By providing a block diagram having a number of systems the method generates a first set of functions and a second set of functions. The first set of functions is associated with a first group of systems that process data at a first rate. The first set of functions invokes the first group of systems implicitly therefrom. The second set of functions is associated with a second group of systems that process data at a second rate. The second set of functions invokes the second group of systems implicitly therefrom. Each of the systems in the first group and the second group are capable of having one or more subsystems that process data at a corresponding rate. In such instances, the method includes a step of generating a set of sub-functions. The set of sub-functions are associated with the sub-systems to invoke the sub-systems implicitly therefrom.

In yet another aspect of the present invention, a method is performed in an electronic device that provides a block diagram environment. A block diagram is provided that includes a block having functions defined in a mechanism outside of the block diagram. Performance of the method identifies portions of the block by a rate of operation and groups code for the block into a number of functions with one function for each portion of the block identified. The step of grouping the code groups the code into non-contiguous sets of code statements, that is, one code group for each function. The mechanism that defines the functions of the block outside of the block diagram environment includes a selected programming environment such as, a structured programming environment or an object oriented programming environment.

In yet another aspect of the present invention a method is performed in an electronic device that provides a block diagram environment. The method includes steps of providing a block that includes two or more components that execute at different rates. The block performs functions defined by a mechanism outside a context of a model in which the block operates. A step of the method separates generated code from the block into two or more sets of code statements with one set of code statements for each execution rate. A further step associates each of the sets of code statements with a corresponding one of the components of the block based on the execution rate of that block. The step of associating each of the sets of code

statements with a corresponding block based on rate performs an implicit association providing an unconditional execution path for each set of code statements.

Brief Description of the Drawings

- 5 Figure 1A depicts a dynamic system described with ordinary differential equations (ODE);
- Figure 1B depicts a dynamic system described with difference equations;
- Figure 1B depicts a dynamic system described with algebraic equations;
- Figure 2 depicts components of a basic block diagram;
- 10 Figure 3 depicts the desired behavior of an integrator block;
- Figure 3 depicts the desired behavior of an integrator block;
- Figure 4 is a flow chart of the sequence of steps used to perform simulation of the block diagram;
- Figure 5 depicts the replacement of a collection of blocks in a block diagram
- 15 with an accumulator block;
- Figure 6A depicts a block diagram and its associated directed graph;
- Figure 6B depicts a linear sorted list generated from the directed graph of Figure 6A;
- Figure 7A depicts an abstract example of a block diagram being executed;
- 20 Figure 7B depicts an abstract view of the execution methods instantiated by the blocks depicted in Figure 7A;
- Figure 7C depicts a sorted list generated from the data dependencies between blocks of Figure 7A;
- Figure 8 depicts a multi-rate system;
- 25 Figure 9 depicts the block diagram of Figure 7A and Figure 8 with associated methods added to the blocks;
- Figure 10 is a flowchart of the sequence of steps followed by a single-tasking execution loop;
- Figure 11A depicts the creation of execution lists from sorted lists in single
- 30 task mode;
- Figure 11B depicts the execution timing of block diagrams in single task mode in timelines synchronized and non-synchronized with real world time;
- Figure 12A depicts the creation of execution lists from sorted lists in multi-task mode;

Figure 12B depicts the execution timing of block diagrams in multi-task mode;

Figure 13 is a flowchart of the overall sequence of steps taken by Simulink in multi-task mode;

5 Figure 14 is a flowchart of the sequence of steps followed by a variable-step solver.

Figure 15 is an exemplary block diagram illustrating an environment suitable for practicing the illustrative embodiment of the present invention.

10 Figure 16 illustrates an exemplary block diagram model suitable for use in practicing the illustrative embodiment of the present invention.

Figure 17 is an exemplary flow chart that provides an overview of steps taken to practice the illustrative embodiment of the present invention.

Figure 18 illustrates exemplary pseudo code for practicing the illustrative embodiment of the present invention.

15

Detailed Description

The illustrative embodiment of the present invention facilitates generation of functions in a block diagram environment. The methods of the illustrative embodiment generate a top level function for each rate used in a block diagram model. Each function contains only blocks executing at a corresponding rate. The methods of the illustrative embodiment of the present invention avoid the need to pass an argument as indicia to invoke a function. The functions generated in accordance with the methods of the illustrative embodiment rely upon implicit identification of the invoking rate from a block in the block diagram model and consequently indicia identifying an operating rate of the block are no longer passed as an argument to the functions. In this manner, functions execute unconditionally based on an execution path free of conditional statements. As such, a block diagram model realizes a reduction in execution time due to an execution path free of conditional statements.

20 The methods of the present invention are equally applicable to blocks in a block diagram environment that includes a functional definition found outside the context of the model in which it is placed. For example, a primitive block found in Simulink®, a software product available from the MathWorks, Inc. of Natick, Massachusetts.

Figure 15 illustrates an environment suitable for practicing the illustrative embodiment of the present invention. A computer system 410 includes an electronic device 412, a network 424, such as the Internet, an intranet, or other suitable network either wired, wireless, or a hybrid of wired and wireless, and, optionally, a server 426 or other electronic device. The electronic device 412 includes a processor 418 for executing various instructions and programs, and controlling various hardware and software components. The electronic device 412 also includes a display device 420 for use in rendering textual and graphical images, a storage device 414 for storing various items such as data, information and programs. A keyboard 422 and a pointing device 424 are also included with the electronic device 412. The pointing device 426 includes such devices as a mouse, track ball, or light pen. Those skilled in the art will recognize that the pointing device 426 can be incorporated with the display device 420 to provide the electronic device 412 with a touch screen that allows a user to interact with the electronic device 412 with a stylus or other means such as a user's finger.

The storage device 414 includes a block diagram environment 416 for use in developing and executing block diagram models. One suitable block diagram environment for use in practicing the illustrative embodiment of the present invention is Simulink[®] from the MathWorks, Inc. of Natick, Massachusetts. Nevertheless, those skilled in the art will appreciate the features of the illustrative embodiment of the present invention described below are equally applicable to other block diagram environments such as LabView, System View, Signal Processing Workstation, HyperSignal, COSSAD, PTOLEMY and other like block diagram environments.

The block diagram environment 416 includes an automatic code generation tool 417. Code generation tool 417 translates a block diagram to executable entities following the layout of the block diagram as provided by the user. The executable entities are compiled and executed on an electronic device such as a computer to implement the functionality specified by the model. Typically, the code generation preserves a model hierarchy in a call graph of the generated code. For instance, each subsystem of a model in a block diagram environment can map to a user specified function and the generated code.

The server 426 coupled to the network 424 includes the block diagram environment 16'. In this manner, a number of users are able to access the block diagram environment 16' via the network 424 without the need to have each user

running a local copy of the block diagram environment 416. Those skilled in the art will recognize that the electronic device 412 includes other software such as, user interfaces and other programs, such as one or more OS programs, compilers and various other program applications developed in a variety of programming environments for controlling system software and hardware components.

The process of developing a block diagram model in the block diagram environment 416 entails a user creating a pictorial model of a dynamic system using blocks. Each block in the model represents a dynamic system whose inputs, outputs, and states change continuously, discretely at points in time or both. If all of a block's inputs, outputs and states change either continuously or at one fixed rate the block is considered single rate. If the inputs, outputs, and states change together or separately at two or more rates the block is considered multi-rate. If a model includes multi-rate blocks or includes two or more single-rate blocks operating at different rates, then the model is considered multi-rate. The model can include blocks that are referred to in the Simulink® block diagram environment as “primitive” or “elementary”. A block is “primitive” if its functional definition is found outside the context of the model in which the block is placed. For example, in the Simulink® block diagram environment predefined primitive blocks include a sum block, a product block, and a gain block. Nevertheless, those skilled in the art will recognize that the Simulink® diagram environment includes other predefined primitive blocks and that a user can create their own primitive blocks in the block diagram environment 416. One manner in which a functional definition of a block is specified outside the context of the block diagram model is for example, using “C” code or MATLAB “M” code developed by MathWorks, Inc. of Natick, Massachusetts. Those skilled in the art will also appreciate that the functionality of a primitive block can be changed, but not from within a block editing tool.

The block diagram environment 416 also includes “atomic sub-systems”. An “atomic sub-system” is a block or group of blocks whose functional definition is specified through the block diagram model in which the block is placed. Often the functional definition of an atomic sub-system is defined using sets of primitive blocks. One manner in which to provide functional definition for an “atomic sub-system” is through the use of interconnected sets of predefined blocks. In this manner, user written functions and other like mechanisms are used to specify an atomic sub-system. An atomic sub-system in the block diagram environment 416 can have a

hierarchical architecture to define a model hierarchy. Those skilled in the art will recognize that in selected instances, a root model of a system is an atomic sub-system.

Figure 16 illustrates a block diagram model 430 suitable for use in practicing the illustrative embodiment of the present invention. The block diagram model 430 includes a first sub-system 432 coupled to a first gain block 436. The first gain block 36 is coupled to a second sub-system 434, which, in turn, is coupled to an integration block 38. The integration block 438 is also coupled to the first sub-system 432.

The first sub-system 432 includes a first block 440 and a fourth gain block 442. The first block 440 receives an output from the integration block 432 and is capable of performing an operation on the output received from the integration block 438 or is alternately capable of passing the received output from the integration block 438 to an input of the fourth gain block 442. The fourth gain block 442 performs an operation to increase an amplitude of the signal received on the input. The first sub-system 32 is a multi-rate system for the block 440 operates at a second operating rate, for example, 10 Hz and the fourth gain block 442 operates at a first operating rate, for example, 5 Hz.

The first gain block 436 performs an amplification operation on the output of the first sub-system 432 and passes the amplified value to the second sub-system 434. The first gain block 436 is a single rate block that operates at the first operating rate.

The second sub-system 434 receives the output of the first gain block 436 at an input of the second gain block 446. The second gain block 446 performs an operation to amplify the received value and passes the amplified value to the third gain block 444. The third gain block 444 in turn, performs an amplification operation on the value received from the second gain block 446 and passes the amplified value to an input of the integration block 438. The second sub-system 434 is a multi rate sub-system for the second gain block 446 operates at the first operating rate and the third gain block 444 operates at the second operating rate.

The integration block 438 integrates the value received on an input from the output of second sub-system 434 and passes the resulting value to an input of the first sub-system 432. The integration block 438 operates at the second operating frequency. The block diagram model 430 is a multi-rate model because it includes two or more blocks or a sub-system that includes blocks or portions of blocks that operate at different rates. Nevertheless, those skilled in the art will recognize that the block diagram model 430 is merely illustrative and not meant to limit the illustrative

embodiment of the present invention. For example, those skilled in the art will recognize that block diagram model 430 can include a significant number of blocks, for example, one hundred or more and a significant number of sub-systems, for example, fifty or more and include more than two operating rates in the block diagram model.

From the block diagram model 432 an automatic code generation tool, for example, Real-Time Workshop from the MathWorks, Inc. of Natick, Massachusetts translates the block diagram model into program code that is compiled and executed on the electronic device 412 to implement the functionality of the block diagram model 430. The generation of the program code from the block diagram model 430 preserves a model hierarchy in a call graph of the generated code. This allows each sub-system in the block diagram model 430 to map to a user specified function in the generated code.

Figure 17 illustrates the steps taken to generate code from a block diagram model 430 in accordance with the illustrative embodiment of the present invention. In Step 350, the code generation tool identifies the various operating rates within the block diagram model 430. For example, a first operating rate and a second operating rate. Those skilled in the art will recognize that a selected block can have more than one operating rate or a sub-system can have more than one block operating at different rates. In Step 352, the code generation tool 417 generates code from the block diagram model 430. The generated code includes one function for each identified operating rate. Each function generated in Step 352 is considered a top level function and contains only blocks executing at the corresponding rate.

In Step 354, sub-systems that request generation of separate functions causes the code generation tool to generate a sub-function for each rate used by the blocks of the requesting sub-system. The sub-functions generated in Step 354 are callable by the top level functions generated in Step 352.

In Step 356, at completion of code generation, execution of the generated code instructs a first function to invoke an operation operating at a corresponding first operating rate. Specifically, instructing a first function to invoke an operation operating at a corresponding first operating rate is an implicit operation. Because the identification of the invoking rate is implicit it is no longer required that indicia identifying the operating rate be passed as an argument to any of the generated functions. Moreover, since the generated code represents only blocks executing at the

rate for which the function is generated there is no longer a need for passing indicia identifying the rate to a function. In Step 358, execution of the generated code instructs the second function to invoke an operation operating at a second operating rate. Again, the function invokes the corresponding operating rate implicitly for the function defines an execution path free of conditional statements for the corresponding rate.

Figure 18 illustrates pseudo code representing code generated in accordance with the illustrative embodiment of the present invention. Based on the block diagram model 430 the code generation tool generates a function 460 for each of the blocks in the block diagram model 430 operating at the first operating rate, for example 5 Hz, and generates a second function 468 for each block in the block diagram model 30 operating at the second operating rate, for example 10 Hz. Function 460 includes a sub-function 462 corresponding to gain block 442 in the first sub-system 432 that executes at the first operating rate and a second sub-function 464 corresponding to gain block 444 in the second sub-system 434 executing at the first operating rate. In like fashion, function 468 is associated with a sub-function 470 corresponding to block 440 in the first sub-system 432 operating at the second operating rate and a second sub-function 472 corresponding to gain block 446 in the second sub-system 434 operating at the second operating rate.

A call to the sub-function 470 invokes execution path 480. Code in execution path 480 corresponds to gain block 440 in the second sub-system 434 operating at the second operating rate. In like fashion, a call to sub-function 472 invoke an execution path 476 invoking block 440 in the first sub-system 432 executing at the second operating rate. In similar fashion, a call to sub-function 462 invokes execution path 474 for executing functions of gain block 442 in the first sub-system 432 operating at the first operating rate. Likewise a call to sub-function 464 invokes execution path 478 that invokes functions of gain block 444 in the second sub-system 434 operating at the first operating rate. Those skilled in the art will recognize that the generated code for such functions and sub-functions include no conditional statements and thus each function and sub-function unconditionally define an execution path for each of the rate in a block diagram model.

Should the code generation tool 417 encounter a “primitive” block or a “primitive” block in the block diagram 430, the code for the primitive block is divided into fragments and distributed into the various functions corresponding to the blocks

rates, for example, function 460 or function 468 illustrated in Figure 18. Like the code generation for atomic sub-systems, the generated code for a primitive block does not include logical predicates nor does the generated code pass arguments to identify the active rate or the rate associated with a portion of the primitive block invoking the function.

5 The generation of code from a block diagram model without conditional statements and thus without conditional execution avoids the need to evaluate sets of logical predicates at run time. The logical predicates consume code memory space and moreover, their run time evaluation reduces throughput of the block diagram model. Consequently, generation of code from the block diagram model free of
10 logical predicates to execute selected functions increases throughput and significantly benefits block diagram models modeling real-time events. Moreover, multiple blocks with the same operating rate are able to call the corresponding function without the need to have a full set of coded functions for all the rates duplicated in memory for
15 each task. As a result, tasks share the generated code in real-time.

While the present invention has been described with reference to an illustrative embodiment thereof, one skilled in the art will appreciate that there are changes in form and detail that may be made without departing from the intended scope of the present invention identified in the pending claims. For example, the
20 above-described code generation technique is applicable to other block diagram environments that perform tasks other than simulating dynamic systems; such other tasks include the control and emulation with various measurement instruments.